arXiv:1612.05019v1 [cs.LO] 15 Dec 2016

An early sign of satisfiability

Eliezer L. Lozinskii
School of Computer Science and Engineering
The Hebrew University, Jerusalem 91904, Israel
email: lozinski@cs.huji.ac.il

Abstract

This note considers checking satisfiability of sets of propositional clauses (SAT instances). It shows that *unipolar sets* of clauses (containing no positive or no negative clauses) provide an *early sign* of satisfiability of SAT instances **before** all the clauses become satisfied in the course of solving SAT problems. At this sign the processing can be terminated by *unipolar set termination, UST* thus before it is usually done by SAT solvers (Table 1). An analysis of benchmark SAT instances used at SAT Competitions shows that UST can speed up solving SAT instances stemming from many real-world problems. The efficiency of UST increases with the *skewnwss* of the SAT set being checked, that is the difference between probabilities of negated and unnegated literals in the set. Many real-world problems, by virtue of their semantics, are skewed (Table 2). The efficiency of UST can be increased by revealing the *hidden skewness* of SAT sets (Table 3).

Keywords: Satisfiability checking, unipolar sets of clauses, early termination, SAT speed -up, skewed benchmarks, hidden skewness.

1 Introduction

SAT is the problem of checking satisfiability of a propositional formula F presented in the conjunctive normal form CNF as a conjunction of a set of *clauses*; each clause is a disjunction of *literals*, each literal is an unnegated or negated propositional *variable*. Let S be the set of clauses of F, and V — the set of variables appearing in S. In the process of checking satisfiability of F each variable $v \in V$ can be assigned a truth value *true* or *false*. Let A be a set of assignments to the variables of V. A clause $C \in S$ is *satisfied* if there is a literal in C assigned *true* in A. F is *satisfiable* if there is a set of assignments that satisfies all the clauses of S, otherwise F is *unsatisfiable*.

Solvers for SAT are based on the famous Davis-Putnam-Logemann-Loveland procedure *DPLL* [8, 7]. In search for a set of assignments A satisfying all the

clauses of S DPLL performs a sequence of *steps*. At each step it assigns a truth value to a variable $v \in V$ not yet assigned, and *applies* this assignment α to S in the following way: for every clause C of S (not yet deleted from S), if C contains a literal assigned *true* by α then C is deleted from S, if C contains a literal L assigned *false* by α then L is deleted from C; if after deleting L from C the latter becomes empty then DPLL encounters a *conflict*, so it *backtracks* cancelling the changes produced by the application of α and then performing another assignment instead of α; otherwise if the application of α causes no conflict then the assignment is *successful* and α is appended to A; DPLL proceeds to the next step; if after deleting C from S the latter becomes empty (meaning that all clauses of S are satisfied by A) then DPLL terminates deciding that S is *satisfiable*; however, if the search is exhausted, but a satisfying set of assignments is not found then S is *unsatisfiable*.

In the course of the fifty-odd years since DPLL was pioneered numerous innovations have been introduced and engineered into the procedure, so that the modern SAT solvers are sophisticated and efficient programs capable of checking satisfiability of very large sets containing millions of clauses (a profound description and analysis of SAT solvers is presented in [1, 10]. Modern solvers use efficient branching strategies for determining next assignment; they perform long backtracking jumps pruning large portions of the search space; when conflicts are encountered, new *learned clauses* are constructed and added to the set to prevent repeated assignments leading to the same conflict; to reduce the amount of book-keeping many solvers using a *lazy* data structure do not update the state of all literals of every clause, but keep track of certain *watched literals*; these solvers determine satisfiability of a set of clauses when all its variables have been assigned successfully [13] (at this stage all clauses of the set are satisfied).

Solvers fully implementing DPLL are *complete* such that for any set of clauses they decide whether it is satisfiable or not [6]. There are solvers that unlike complete ones do not carry out an exhaustive search of a satisfying assignment, but perform a stochastic *local search* with no guarantee of finding an existing one [12]. Although the solvers based on local search often outperform complete solvers, they are *incomplete* since for unsatisfiable sets and even for some satisfiable ones these solvers terminate undecided.

2 Unipolar sets

Consider a set of clauses S that is being checked for satisfiability by a SAT solver. In the course of this process the contents and size of S change. At any step when an assignment α is applied to S, clauses satisfied by α are deleted from S or marked as *inactive* (depending on the working solver), so the set of *active* clauses $\tilde{S}$ subject to the further processing shrinks. On the other hand, when a conflict (an empty clause) is encountered and so the search backtracks, some

inactive or previously deleted clauses may be returned into $\tilde{S}$ extending it.

Such a "pulsating" behaviour is characteristic of any set of clauses processed by any of a wide variety of solvers. And important: no matter what is the method of making assignments, backtracking or book-keeping implemented by a solver — when the solver terminates reporting satisfiability of a set S, then there is an assignment found by the solver that satisfies all the clauses of S (although solvers using a lazy data structure [13] may be unaware of this fact). Let us call this event the *all-satisfied termination, AST.*

Definition 2.1 *If a clause C contains unnegated literals only then call C a* positive clause*; if C contains negated literals only then C is a* negative clause*; otherwise C is a* mixed clause*. If a set S contains both a positive and a negative clauses then call S a* bipolar set*; otherwise S is an* unipolar set. □

Observation 2.1 *Any unipolar set S of clauses is satisfiable. Indeed, if S contains no positive (negative) clauses then every clause of S contains a negated (unnegated) literal. So the assignment of* false *(*true*) to all variables of S satisfies all clauses of S.* □

By Observation 2.1, the unipolarity of a set of clauses is a sign of its satisfiability, so if in the course of checking satisfiability of a set S the set of active clauses $\tilde{S}$ becomes unipolar, then at this moment the search can be terminated deciding that S is satisfiable. Let us call this event the *unipolar set termination, UST.* When UST occurs, $\tilde{S}$ is not empty containing clauses that are not yet satisfied by the previous assignments, but are known to be satisfiable by assigning *true* or *false* to all yet unassigned variables of the set. As the steps of variable assignments are performed in a sequence, let α_i be the truth value assignment to a variable at step i. Let $\tilde{S}_{bi}$ and $\tilde{S}_{ai}$ denote, respectively, the states of $\tilde{S}$ immediately before and after application of α_i.

Proposition 2.1 *For all satisfiable sets S of clauses and all sequences of variable assignments, $\tilde{S}$ becomes unipolar* before *all clauses of S are satisfied, so UST is* always achieved in less steps than AST.

Proof. Let α_t be the assignment of *true* to a literal L such that all clauses of $\tilde{S}_{bt}$ become satisfied. So S is satisfiable, and AST takes place after performing t steps of variable assignment. Consider the content of $\tilde{S}_{bt}$ just before the application of α_t: $\tilde{S}_{bt}$ is not empty (otherwise AST would have occurred before step t); all clauses of $\tilde{S}_{bt}$ are satisfied by α_t so every clause of $\tilde{S}_{bt}$ contains L; but a positive and a negative clauses cannot contain the same literal, so $\tilde{S}_{bt}$ is unipolar: either no positive or no negative clauses. Should $\tilde{S}_{bt}$ be checked for unipolarity, satisfiability of S would be detected, and UST performed before step t, so step t would not be needed. □

Unipolarity is a common feature of satisfiable sets of clauses in the following sense. Let θ denote a subset of variables appearing in a set of clauses S, and $S\theta$ stand for a set resulting from inverting all literals in S involving the variables in θ. Call θ an *inverter*.

Proposition 2.2 *A set of clauses S is satisfiable iff there exists an inverter θ such that $S\theta$ is unipolar.*

Proof. *If:* For all sets S and all inverters θ, S is satisfiable iff $S\theta$ is so. Indeed, if M is a model of S then $M\theta$ is a model of $S\theta$, and vice versa. So if $S\theta$ is unipolar then, by Observation 2.1, S is satisfiable.

Only if: If S is satisfiable, and M is one of its models, define an inverter θ as the set of all unnegated variables in M. $M\theta$ contains negated literals only and is a model of $S\theta$, hence, $S\theta$ contains no positive clause and so is unipolar. □

3 The gain of UST over AST

To evaluate the efficiency of the UST relative to the AST we have run experiments with a program that implements DPLL: given a set of caluses S, the program at each step assigns *true* to the most frequent literal of a most frequent variable among the active clauses; after applying every variable assignment the program updates the size of the sets of active positive and negative clauses, and if one of them is empty then at this step the processing could be terminate by UST, but the search goes on until all clauses of the set are satisfied, so AST is performed. Let N_U and N_A denote the number of variable assignments made by the program till UST and AST, respectively, are reached, and $G = N_A/N_U$ stand for the *gain* of UST over AST. By Proposition 2.1, $N_U < N_A$, so $G > 1$. Let R denote the *remainder*, that is the percentage of the clauses of S remained active but not yet satisfied at UST; $R > 0$. The larger the values of G and R, the more efficient UST is for S. The sets S in the experiments were generated with the following parameters (Table 1): number of variables $n = 100$; the clauses-to-variables ratio $r = m/n$ varied from $r = 2$ to the threshold value (shown in boldface in Table 1) at which satisfiability of S undergoes phase transition [5]; 3 literals in every clause; all variables appear in S with the same probability, however negated and unnegated literals have different probabilities; without loss of generality, the probability of an unnegated literal $p \leq 0.5$; 1000 instances were checked for each pair of values (p, r).

Although sets for testing SAT solvers often are generated with $p = 0.5$, sets with $p < 0.5$ are theoretically interesting and practically important (next section). Sets with different probability of negated and unnegated literals (*skewed* sets) were studied by Sinopalnikov [18], and shown to undergo satisfiability phase transition at a threshold value of r that grows with decreasing value of p. The efficiency of UST depends on the value of p. Indeed, the smaller the value of p in

Table 1: The gain G and remainder R for $p = 0.5 - 0.05$, $n = 100$, 1000 tests for each pair (p, r)

p	r	2.00	2.25	2.50	2.75	3.00	3.25	3.50	3.75	4.00	**4.26**
0.5	G	1.08	1.07	1.07	1.06	1.05	1.02	1.01	1.00	1.00	1.00
	$R\%$	3	3	3	2	2	2	2	1	1	1
p	r	2.00	2.30	2.60	2.90	3.20	3.50	3.80	4.10	4.40	**4.70**
0.4	G	1.11	1.09	1.08	1.07	1.06	1.03	1.01	1.00	1.00	1.00
	$R\%$	4	3	3	3	2	2	2	2	1	1
p	r	2.00	2.50	3.00	3.50	4.00	4.50	5.00	5.50	6.00	**6.40**
0.3	G	1.22	1.18	1.15	1.12	1.10	1.05	1.02	1.00	1.00	1.00
	$R\%$	8	6	5	4	3	3	2	2	2	1
p	r	2.00	3.00	4.00	5.00	6.00	7.00	8.00	9.00	10.00	**11.5**
0.2	G	1.85	1.54	1.38	1.30	1.23	1.15	1.07	1.02	1.01	1.01
	$R\%$	33	21	13	9	7	6	4	3	3	2
p	r	2.00	3.00	5.00	10.0	15.0	20.0	25.0	30.0	35.0	**41.0**
0.1	G	8.60	6.36	4.33	2.44	1.89	1.58	1.26	1.07	1.06	1.06
	$R\%$	76	72	69	45	31	25	16	12	9	9
p	r	2.00	10.0	30.0	50.0	70.0	90.0	110	130	150	**165**
0.05	G	48.8	14.8	5.47	3.28	2.24	1.46	1.32	1.31	1.31	1.30
	$R\%$	98	91	74	58	48	41	34	33	33	32

S, the smaller the probability that a positive clause appears in S and so larger the probability that S becomes unipolar much earlier than all its clauses become satisfied in the process of checking its satisfiability.

Results of the experiments are summarised in Table 1. While for $p \geq 0.3$ UST gains just several percents of assignments over AST, for $p < 0.3$ UST requires many times less assignments than AST. For instance, for $n = 100, p = 0.1, r = 5.0$, UST is 4.33 times faster than AST: UST detects satisfiability after 15 assignments; at this step there remain 345 active clauses not yet satisfied (out of 500 initially), and then AST requires 50 more assignments (averaged over 1000 instances). It is remarkable that satisfiability of a set can be decided when a significant part of its clauses (69% in this example) have not been satisfied.

4 "Skewness" of real-world SAT problems

SAT is important not only in the theory of computation being the core NP-complete problem [4], but it has numerous practical applications, since many real-world problems can be encoded as SAT instances. Among these SAT en-

Table 2: Skewness parameters of practical SAT instances

Source	**Family**	**# inst**	**p**
SATLIB [16]	Planning blocksworld	7	0.14 - 0.26
	Planning logistics	4	0.28 - 0.42
	Graph colouring	400	0.08 - 0.13
	Quasigroup encoded	22	0.003 - 0.278
SAT-2005 [20]	Maris	67	0.03 - 0.33
	Grieu	12	0.24
SAT-2014 [14]	Oldpool	8	0.15 - 0.40
	Wallner	20	0.07 - 0.15

codings are such important problems and techniques as planning, digital circuits design, software verification, network design and analysis, diagnosing, data security, cryptanalysis, graph colouring, proof checking, automated reasoning [2, 3, 9, 11, 19, 21, 22].

SAT encodings stemming from practical problems are likely to be "skewed", since unnegated and negated literals represent different aspects of the encoded problems. For example, unnegated literals usually represent certain features of the corresponding objects, while negated literals participate in encoding of restrictions and conditions imposed upon these features. The restrictions involve usually combinations of features and are often more numerous than individual features and objects. So by virtue of their semantics many practical problems produce skewed SAT instances.

Table 2 shows skewness parameters of several benchmark SAT instances of the category *industrial* and *applications* from SATLIB [16] and SAT Competitions [14, 20] (540 instances). For almost all these instances, their values of $\mathbf{p}$ fall within the range of a significant gain of UST (Table 1). These findings suggest that performing UST *instead of* AST would speed up practical SAT solvers. Notably, all 13 instances of the subset */wallner-argumentation/complete* from SAT 2014 competition (part of the last line of Table 2) are unipolar initially (no positive clauses), so by UST the processing would be terminated at the very first step "in no time". This fact went by unnoticed such that these instances appeared again as benchmarks at the SAT Race 2015 Competition (www.baldur.iti.kit.edu/sat-race-2015/index.php?cat=downloads).

Although UST requires less assignments than AST, it involves more book-keeping for updating the numbers of positive and negative active clauses. The efficiency of UST depends strongly on the skewness of the set. In order to decide whether to perform UST for a given set (or spare this additional book-keeping for a set with the value of $\mathbf{p}$ close to 0.5) the value of $\mathbf{p}$ can be counted in linear runtime at a preprocessing. For significantly skewed SAT sets, UST is much faster

than AST, so for this kind of SAT instances corresponding to many practical problems, all solvers, complete and incomplete, can benefit from performing UST.

5 Hidden skewness

The initial skewness of a SAT instance can be enhanced without affecting its satisfiability, so increasing the efficiency of UST. Let $pos(v, S)$, $neg(v, S)$ denote, respectively, the number of unnegated, negated occurrences of a variable v in a set of clauses S, and $poslit(S)$, $neglit(S)$ stend for the total number of unnegated, negated literals in S. Then the skewness $p(S)$ of S is

$$p(S) = min(poslit(S), neglit(S))/(poslit(S) + neglit(S)).$$

Definition 5.1 *Given a set S, define an inverter $\rho_S = \{v | pos(v, S) > neg(v, S)\}$. We say that ρ_S* reveals *the* hidden skewness *of S, $hp(S)$, such that $hp(S) = p(S\rho_S)$* $\square$

Proposition 5.1 *For all sets S, $hp(S) \leq p(S)$.*

Proof. By inverting in S all literals involving the variables of ρ_S we get:
for all $v \in \rho_S$,
$pos(v, S\rho_S) < neg(v, S\rho_S)$, $pos(v, S\rho_S) < pos(v, S)$, $pos(v, S\rho_S) = neg(v, S)$,
while for all $v \notin \rho_S$,
$pos(v, S\rho_S) \leq neg(v, S\rho_S)$, $pos(v, S\rho_S) = pos(v, S)$, $pos(v, S\rho_S) \leq neg(v, S)$.
Hence,
$poslit(S\rho_S) \leq neglit(S\rho_S)$, $poslit(S\rho_S) \leq poslit(S)$, $poslit(S\rho_S) \leq neglit(S)$,
and so $p(S\rho_S) \leq p(S)$. $\square$

Table 3 shows hidden skewness of a sample of benchmark instances from the SAT Race 2015 and SAT 2016 Competitions (www.baldur.iti.kit.edu/(sat-race-2015 and sat-competition-2016)/index.php?cat=downloads). For many SAT sets $hp(S)$ is significantly less than $p(S)$, so revealing the hidden skewness of a given set increases the efficiency of UST. Inversion of literals involving the variables of δ can be performed at a preprocessing in a linear run-time.

Example 5.1 *Given a set $S = \{(v_1, v_2, v_3), (\neg v_1, \neg v_2, \neg v_3), (\neg v_1, v_2, \neg v_3)\}$, we get: $p(S) = 0.44, \rho_S = \{v_2\}$, $S\rho_S = \{(v_1, \neg v_2, v_3), (\neg v_1, v_2, \neg v_3), (\neg v_1, \neg v_2, \neg v_3)\}$ and the hidden skewness of S, $hp(S) = p(S\rho_S) = 0.33$. By the way, $S\rho_S$ is already unipolar.* $\square$

6 Conclusion

The *unipolar set termination*, *UST*, is presented that detects satisfiability in the process of solving SAT *always before* the all-satisfied termination, AST, performed usually by SAT solvers, takes place.

Table 3: Hidden vs initial skewness of benchmark SAT instances

Source	Instance (File) S	#var n	#cla m	Initial p(S)	Hidden hp(S)
SAT-Race-2015 [17]	jgiraldezlevi.2200.9086.08.40.108.cnf	2200	9086	0.499	0.388
	manthey_single-ordered-initialized-w18-b8.cnf	2160	15054	0.407	0.364
	partial-5-17-s.cnf	252328	1189896	0.407	0.359
	aaai10-planning-ipc5-TPP-21-step11.cnf	99736	783991	0.352	0.260
SAT-2016 [15]	C168_FW_UT_518.cnf	1909	7511	0.446	0.234
	gripper14u.cnf	4584	43390	0.249	0.172
	korf-18.cnf	7794	186934	0.169	0.027
	E00N23.cnf	15364	2210893	0.150	0.004

We measure the efficiency of UST by its *gain*, G, that is the ratio of the number of variable assignments required by AST to that of UST. Table 1 shows values of G produced by experiments with SAT instances generated with varied parameters. UST is most efficient for *skewed* sets of clauses with different probability of unnegated and negated literals. It turns out that SAT instances encoding real-world problems are very likely to be significantly skewed. This is true of many instances stemming from practical problems used as benchmarks for SAT competitions (shown in Tables 2, 3). Implementing UST requires updating the number of positive and negative active clauses at every variable assignment, but as suggested by this short study, for many important practical SAT problems the gain of UST over AST can by far compensate for this additional book-keeping, thus speeding up the process of solving SAT. Revealing the hidden skewness of a SAT sets (Table 3) can increase the efficiency of UST.

References

[1] Biere, A., Heule, M., van Maaren, H. and Walsh, T., Eds. (2009) *Handbook of Satisfiability*, IOS Press.

[2] Bryant, R. E., German, S. M. and Velev, M. N. (1999) Microprocessor Verification Using Efficient Decision Procedures for a Logic of Equality with Uninterpreted Functions, in *Analytic Tableaux and Related Methods*, 1–13.

[3] Clarke, E., Biere, A., Raimi, R. and Zhu, Y (2001) Bounded Model Checking Using Satisfiability Solving, *Formal Methods in System Design*, 19: 7.

[4] Cook, S. (1971) The complexity of theorem proving procedures, *Proceedings of 3rd ACM STOC*, 151–158.

[5] Crawford, J. and Auton, L. (1996) Experimental results on the crossover point in random 3-SAT, *Artificial Intelligence*, **57**: 31–81.

[6] Darwiche, A. and Pipatsrisawat, K. (2009) Complete algorithms, *Handbook of Satisfiability*, Armin Biere, Marijn Heule, Hans van Maaren and Toby Walsh, Eds., 99–130, IOS Press.

[7] Davis, M., Logemann, G. and Loveland, D. (1962) A machine program for theorem proving, *Communications of the ACM*, **5** (7): 394–397.

[8] Davis, M. and Putnam, H. (1960) A computing procedure for quantification theory, *Journal of the ACM*, **7**: 201–215.

[9] Een, N. and Sorensson, N. (2006) Translating pseudo-Boolean constraints into SAT. *Journal on Satisfiability, Boolean Modeling and Computation*, 2:1–26.

[10] Gomes, C. P., Kautz, H., Sabharwal, A. and Selman, B. (2008) Satisfiability solvers, *Handbook of Knowledge Representation*, F. van Harmelen, V. Lifschitz and B. Porter, Eds., 89–134, Elsevier.

[11] Kautz H. A. and Selman, B. (1996) Pushing the envelope: Planning, propositional logic, and stochastic search. In *13th AAAI*, 1194–1201, Portland, OR.

[12] Kautz, H., Sabharwal, A. and Selman, B. (2009) Incomplete algorithms, *Handbook of Satisfiability*, Armin Biere, Marijn Heule, Hans van Maaren and Toby Walsh, Eds., 185–204, IOS Press.

[13] Marques-Silva, J., Lynce, I. and Malik, S. (2009) Conflict-driven clause learning SAT solvers, *Handbook of Satisfiability*, Armin Biere, Marijn Heule, Hans van Maaren and Toby Walsh, Eds., 131–154, IOS Press.

[14] SAT Competition 2014:
www.satcompetition.org/edacc/sc14/experiment/18/instances

[15] SAT Competition 2016:
www.baldur.iti.kit.edu/sat-competition-2016/index.php?cat=downloads

[16] SATLIB Benchmark Problems:
www.cs.ubc.ca/∼hoos/SATLIB/benchm.html

[17] SAT Race 2015:
www.baldur.iti.kit.edu/sat-race-2015/index.php?cat=downloads

[18] Sinopalnikov, D. A. (2004) Satisfiability threshold on the skewed random k-SAT, *Proceedings, 7th Intl. Conf. on Theory and Applications of Satisfiability Testing*, Vancouver, Canada.

[19] Stephan, P. R., Brayton, R. K. and Sangiovanni-Vincentelli, A. L. (1996) Combinatorial test generation using satisfiability. *IEEE Transactions on CAD and IC*, 15(9):1167–1176.

[20] The International SAT Competitions web page: *www.satcompetition.org*

[21] Velev, M. N. and R. E. Bryant, R. E. (2003) Effective use of Boolean satisfiability procedures in the formal verification of superscalar and vliw microprocessors. *J. Symb. Comput.*, 35(2):73–106.

[22] Vizel, Y., Weissenbacher, G. and Malik, S. (2015) Boolean Satisfiability Solvers and Their Applications in Model Checking, *Proceedings of the IEEE*, 103 (11).